Computing in Memory with Spin-Transfer Torque Magnetic RAM

Shubham Jain, Ashish Ranjan, Kaushik Roy, Anand Raghunathan
School of Electrical and Computer Engineering, Purdue University
{jain130,aranjan,kaushik,raghunathan}@purdue.edu

arXiv:1703.02118v3 [cs.ET] 4 Apr 2017

***Abstract*—In-memory computing is a promising approach to reducing the time and energy spent on data transfers between the processor and memory, thereby alleviating the processor-memory gap. We explore in-memory computing with STT-MRAM, which is considered a promising candidate for future on-chip memories due to its non-volatility, density, and near-zero leakage. We propose suitable modifications to peripheral circuits that enable standard STT-MRAM arrays to perform bitwise, arithmetic and complex vector operations. We address the key challenge of computing reliably under process variations, by leveraging ECC schemes that are employed for conventional memory operations to also correct errors during in-memory computations. We propose architectural enhancements to the instruction set and on-chip bus that enable the proposed design, Spin-Transfer Torque Compute-in-Memory (STT-CiM), to be integrated into a programmable computing system. We also present data mapping techniques to increase the effectiveness of STT-CiM. We evaluate STT-CiM using a device-to-architecture modeling framework, and integrate cycle-accurate models of STT-CiM with a commercial processor and on-chip bus (Nios II and Avalon from Intel). Our system-level evaluation shows that STT-CiM provides system performance improvements of 3.93X on average (upto 12.4X), and concurrently reduces memory system energy by 3.83X on average (upto 12.4X).**

I. Introduction

The growth in data sets and increase in the number of cores place high demands on the memory systems of modern computing platforms. Consequently, a growing fraction of transistors, area and power are utilized towards memories. CMOS memories (SRAM and embedded DRAM) have been the mainstays of memory design for the past several decades. However, recent technology scaling challenges in CMOS memories, along with an increased demand for memory capacity and performance, have fueled an active interest in alternative memory technologies.

Spintronic memories have emerged as a promising candidate for future memories due to several desirable attributes such as non-volatility, high density, and near-zero leakage. In particular, Spin Transfer Torque Magnetic RAM (STT-MRAM) has garnered significant interest with various prototype demonstrations and early commercial offerings [1]–[3]. There have been several research efforts to boost the efficiency of STT-MRAM at the device, circuit and architectural levels [4]–[26]. In this work, we explore a complementary direction, *viz* in-memory computing [27]–[33], by enhancing STT-MRAM arrays to perform a range of arithmetic, logic and vector operations. We propose circuit and architectural techniques to address the challenges associated with designing and using such structures, including reliable computation under process variations and architectural support to enable the structures to be used in a programmable processor based system.

In-memory computing is motivated by the observation that the movement of data from bit-cells in the memory to the processor and back (across the bit-lines, memory interface, and system interconnect) is a major bottleneck to performance and energy efficiency of computing systems [34].

Efforts that have explored the closer integration of logic and memory are variedly referred to in the literature (logic-in-memory, compute-in-memory, processing-in-memory, *etc.*), and may be classified into three categories – moving logic closer to memory [35]–[47], performing computations within memory structures [27]–[33], and embedding nonvolatile storage elements within logic [48]–[51]. The first two approaches address the efficiency of performing active computation, whereas the third addresses the challenge of memory energy during idle periods.

Our proposal addresses the problem of in-memory computing with spintronic memories, to improve performance and energy efficiency without adversely impacting density or efficiency as a conventional memory. We are aware of only one other effort on in-memory computing with spintronic memory [29], which requires the addition of a transistor to each bit-cell (resulting in a 2T-1R cell) [29]. Our proposal differs significantly in that it *requires no changes to the bit-cells and core array*, enabling it to be used without sacrificing memory density. Our proposal is based on the observation that by enabling multiple wordlines simultaneously [1], and sensing the effective resistance of all the enabled bit-cells in each bit-line, it is possible to directly compute logic functions of the values stored in the bit-cells. We note that Pinatubo [30] first proposed enabling multiple wordlines to perform computations within Non-Volatile Memories (NVMs). Although our work shares this mechanism with Pinatubo, we differ from (and go beyond) it in several key aspects. First, we address the key challenge of reliable in-memory computing under process variations. While reliable sensing under the limited tunneling magneto-resistance (TMR) of STT-MRAM bit-cells is known to be a challenge [13]–[16], this is unfortunately further aggravated by in-memory computations. We show how to leverage

[1]Note that this will lead to short circuit paths in SRAMs, but not in STT-MRAM due to the resistive nature of bit-cells.

conventional ECC schemes to also improve the reliability of in-memory computations. Second, we propose peripheral modifications that enables us to go beyond bitwise operations by also performing arithmetic and vector operations. Finally, unlike Pinatubo that focuses on off-chip memories, we explore architectural enhancements (bus and ISA extensions), and data mapping techniques to enable in-memory computation in the context of on-chip memories.

In this work, we propose STT-CiM, an STT-MRAM memory array that is enhanced to perform a range of arithmetic, logic, and vector operations. In STT-CiM, the core data array is the same as standard STT-MRAM; hence, density and efficiency of read and write operations are maintained. In order to make STT-CiM robust under process variations, we extend error correction codes (ECC) to be used for errors that occur during in-memory computations. To evaluate the benefits of STT-CiM, we utilize it as a scratchpad in the memory hierarchy of the Intel Nios II [52] processor. We propose enhancements to the on-chip bus and extend the instruction set of the processor to support compute-in-memory operations and expose them to software. We also present suitable data mapping techniques to maximize the benefits of STT-CiM.

In summary, the key contributions of this work are as follows:

- We explore compute-in-memory with spintronic memories as an approach to improving system performance and energy.
- We propose STT-CiM, an enhanced STT-MRAM array that can perform a range of arithmetic, logic and vector compute-in-memory operations without modifying either the bit-cells or the core data array.
- We address a key challenge in STT-CiM, *i.e.* reliably performing in-memory operations under process variation, by demonstrating suitable error correction mechanisms.
- We propose extensions to the instruction set and on-chip bus to integrate STT-CiM into a programmable processor system and demonstrate the viability of these extensions using Intel's Nios II processor and Avalon on-chip bus.
- We evaluate the performance and energy benefits of STT-CiM, achieving average improvements of 3.83x (upto 12.4X) and 3.93x (upto 10.4X) in the total memory energy and system performance, respectively.

The rest of the paper is organized as follows. Section II presents an overview of prior research efforts related to in-memory computation. Section III provides the necessary background on STT-MRAM. Section IV describes the proposed STT-CiM design and how it supports compute-in-memory operations. Section V outlines the architectural optimizations and support for STT-CiM. Section VI describes the experimental methodology and the application benchmarks. Experimental results are presented in section VII. Section VIII concludes the paper.

II. Related Work

Related efforts can be broadly classified into three categories, as shown in Figure 1. We note that subsets of these efforts are variedly referred to in the literature as logic-in-memory, processing-in-memory and computing-in-memory.

Near-Memory Computing (Placing logic near memory)		**Memory in Logic** (Distributing small memory elements over logic plane)	**In-Memory Computing** (Moving computations within memory array itself)	
General-purpose e.g. 3D stacking (HMC, HBM) [46,47], IRAM [35]	Application-specific e.g. Interpolation memory [45]	e.g. Non-volatile processors, Non-volatile compute units (adders, multipliers, etc.) [48-50]	Application-specific e.g. TCAMs [27], Compute-Memory [28]	General-purpose e.g. STT-CiM Pinatubo [30], bulk bitwise in-DRAM [31] STT-CiM Focus

Fig. 1: Related work: Overview

Near-memory computing refers to placing logic or processing units next to memory. It has been explored at various levels of the memory hierarchy [35]–[47]. Intelligent RAM (IRAM) [35] was an early effort that integrated processor and DRAM in the same chip to improve bandwidth. Embedding simple processing units within each page of main memory [36] and within secondary storage [37] enables computations to be performed near memory. An application-specific example of near-memory computation is memory that can generate interpolated values, enabling the evaluation of complex mathematical functions [45]. Near-memory computing has gained significant interest in recent years, with industry efforts like Hybrid Memory Cube (HMC) [46] and High Bandwidth Memory (HBM) [47].

Memory-in-logic involves embedding non-volatile memory elements into a logic circuit [48]–[51]. The objective is typically to provide non-volatility, enabling the system to efficiently transition from shut-down to active state.

In-memory computing [27]–[33] fundamentally differs from the above approaches since it integrates logic operations into the memory array itself, thereby achieving higher efficiency. Moreover, it is also complementary to the aforementioned efforts, since it does not preclude additional computations being performed near memory or non-volatile storage elements being embedded into the logic. The key challenge of in-memory computing is to realize it without impacting the desirability of the resulting design as a standard memory (*i.e.*, density or efficiency of standard read and write operations).

We can classify previous efforts on in-memory computing based on whether they target application-specific or general-purpose computes, and based on the underlying memory technology that they consider. Application-specific examples of in-memory computing include sum-of-absolute difference [28] and dot-product [33] computation. Ternary content-addressable memory [27], ROM-embedded RAM [53], AC-DIMM [32] and Micron's automata processor [54] can also be viewed as examples of in-memory computing that target specific operations such as pattern matching.

Unlike these application-specific designs, we focus on embedding a broader class of operations (arithmetic, logic and vector operations) within memory. Pinatubo [30] and bulk bitwise in-DRAM [31] explored bitwise in-memory computations for off-chip main memory using Phase Change Memory (PCM) and DRAM, respectively. Our work differs from these efforts in several important aspects. First, we focus on in-memory computing for on-chip spintronic memory that

involves a fundamentally different design space. We propose architectural extensions (bus interface support and ISA extension) and data mapping techniques to enable in-memory computing within a general-purpose processor system. Second, our design is not restricted to bitwise operations and realizes arithmetic as well as complex vector operations. As compared to Pinatubo [30], we use a different sensing circuitry (consisting of 2 sense amplifiers) and reference generation circuitry, which enables us to realize a wider variety of operations. Third, the proposed design requires only one array access to perform bit-wise XOR operations (unlike two in the case of [30]) and the proposed operations are not destructive to the contents stored in the accessed bit-cells (unlike [31]). Finally, we address a key challenge associated with in-memory computations performed by enabling multiple wordlines in NVMs, *i.e.*, performing these operations reliably under process variations. One other recent effort that addresses computing with spintronic memories [29] uses an extra transistor in each bit-cell (2T-1R cells), which sacrifices the density benefits of standard (1T-1R) STT-MRAM. In contrast, our proposal (STT-CiM) enables in-memory computation within a standard STT-MRAM array with no changes to the bit-cells. We also propose architectural support for STT-CiM and evaluate its performance and energy benefits in the context of a processor-based system.

III. Background

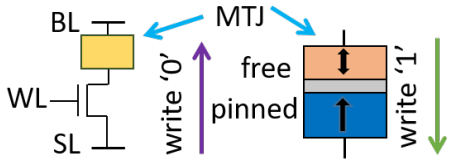

Fig. 2: STT-MRAM bit-cell

An STT-MRAM bit-cell consists of an access transistor and a magnetic tunnel junction (MTJ), as shown in Figure 2. An MTJ in turn consists of a pinned layer that has a fixed magnetic orientation and a free layer whose magnetic orientation can be switched, separated by a tunnelling oxide. The relative magnetic orientation of the free and pinned layers determines the resistance offered by the MTJ (the resistance for the parallel configuration, R_P, is lower than the anti-parallel resistance, R_{AP}). The two resistance states encode a bit (we assume that parallel represents logic "1", and anti-parallel represents logic "0"). A read operation is performed by applying a bias (V_{read}) between the bitline (BL) and the source line (SL), and enabling the wordline (WL). The resultant current flowing through the bit-cell (I_P or I_{AP}) is sensed against a global reference to determine the logic state stored in the bit-cell. A write is performed by passing a current greater than the critical switching current of the MTJ. The logic value written in the bit-cell is dependent on the direction of the write current as shown in Figure 2. The write operation in STT-MRAM is stochastic in nature, and the duration and magnitude of the write current determines the write failure rate. Apart from write failures, STT-MRAM may also be subject to read decision failures, where the value stored in a bit-cell is incorrectly read due to process variations, and read disturb failures where a read operation inadvertently ends up writing into the cell. These failures are addressed through a range of techniques including device and circuit optimization, manufacturing test and self-repair, and error correcting codes [13]–[16].

IV. STT-MRAM based Compute-in-Memory (STT-CiM)

In this section, we describe STT-MRAM based Compute-in-Memory (STT-CiM), which enhances a standard STT-MRAM array to enable in-memory computation.

A. STT-CiM overview

The key idea behind STT-CiM is to enable multiple wordlines simultaneously in an STT-MRAM array, leading to multiple bit-cells being connected to each bitline. With enhancements that we propose to the sensing and reference generation circuitry, we can directly compute logic functions of the enabled words. Note that such an operation is feasible in STT-MRAMs since the bit-cells are resistive. In contrast, enabling multiple wordlines in SRAM can lead to short-circuit paths through the memory array, leading to loss of data stored in the bit-cells.

Figure 3 explains the principle of operation in STT-CiM. First, consider the resistive equivalent circuit of a single STT-MRAM bit-cell shown in Figure 3(a). R_t represents the on-resistance of the access transistor and R_i the resistance of the MTJ. When a voltage bias (V_{read}) is applied between the bitline (BL) and the source line (SL), the net current (I_i) flowing through the bit-cell can take two possible values depending on the MTJ configuration, as shown in Figure 3(b). A read operation involves using a sensing mechanism to distinguish between these two current values.

Figure 3(c) demonstrates a Compute-in-Memory (CiM) operation, where two wordlines (WL_i and WL_j) are enabled, and a voltage bias (V_{read}) is applied to the bitline. The resultant current flowing through the SL (denoted I_{SL}) is a summation of the currents flowing through each of the bit-cells (I_i and I_j), which in turn depends on the logic states stored in these bit-cells. The possible values of I_{SL} are shown in Figure 3(d). We propose enhanced sensing mechanisms to distinguish between these values and thereby compute logic functions of the values stored in the enabled bit-cells. We discuss the details of these operations in turn below.

Bitwise OR (NOR). In order to realize logic OR and NOR operations, we use the sensing scheme shown in Figure 4(a), where I_{SL} is connected to the positive input of the sense amplifier and a reference current I_{ref-or} is fed to its negative input. We choose I_{ref-or} to be between I_{AP-AP} and I_{AP-P}, as shown in Figure 4(c). As a result, among the possible values of I_{SL} [Figure 3(d)], only $I_{SL} = I_{AP-AP}$ is less than I_{ref-or}. Consequently, only the case where both bit-cells are in the AP configuration, *i.e.*, both store "0", leads to an output of logic "0" ("1") at the positive (negative) output of the

sense amplifier, while all other cases lead to logic "1" ("0"). Thus, the positive and negative outputs of the sense amplifier evaluate the logic OR and NOR of the values stored in the enabled bit-cells.

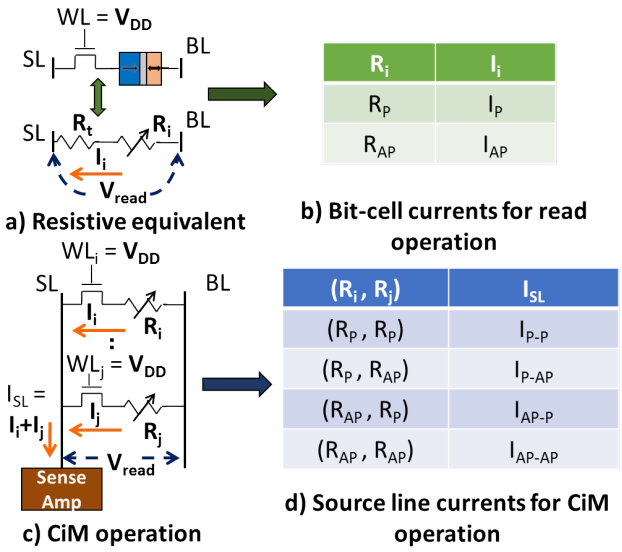

R_i	I_i
R_P	I_P
R_{AP}	I_{AP}

(R_i, R_j)	I_{SL}
(R_P, R_P)	$I_{P\text{-}P}$
(R_P, R_{AP})	$I_{P\text{-}AP}$
(R_{AP}, R_P)	$I_{AP\text{-}P}$
(R_{AP}, R_{AP})	$I_{AP\text{-}AP}$

Fig. 3: STT-CiM: Principle of operation

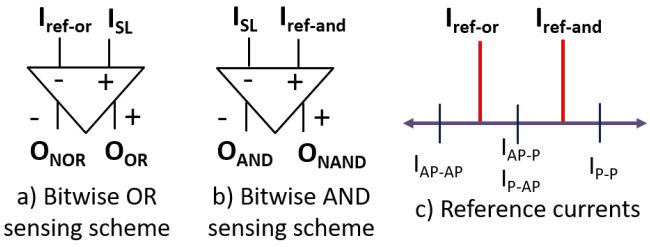

Fig. 4: STT-CiM sensing schemes

Bitwise AND (NAND). A bitwise NAND (AND) operation is realized at the positive (negative) terminal of the sense amplifier by using the sensing scheme shown in Figure 4(b). Note that in this scheme, the reference current ($I_{ref-and}$) is fed to the positive terminal of the sense amplifier.

Bitwise XOR. A bitwise XOR operation is realized when the two sensing schemes shown in Figure 4 are used in tandem, and O_{AND} and O_{NOR} are fed to a CMOS NOR gate. In other words, O_{XOR} = O_{AND} NOR O_{NOR}.

TABLE I: Possible outputs of various sensing schemes

I_{SL}	O_{OR}	O_{NOR}	O_{AND}	O_{NAND}	O_{XOR}
I_{AP-AP}	0	1	0	1	0
I_{AP-P}	1	0	0	1	1
I_{P-AP}	1	0	0	1	1
I_{P-P}	1	0	1	0	0

Table I summarizes the logic operations achieved using the two sensing schemes discussed above. Note that, all the logic operations described above are symmetric in nature, and hence it is not necessary to distinguish between the cases where the two bit-cells connected to a bitline store "10" and "01".

ADD Operation. An ADD operation is realized by leveraging the bitwise logical operations, as illustrated in Figure 5. Suppose A_n and B_n (the n-th bits of two words, A and B) are stored in two different bit-cells of the same column within an STT-CiM array. Suppose that we wish to compute the full-adder logic function (n-th stage of an adder that adds words A and B). As shown in Figure 5, A_n and B_n are not required individually; rather, knowing A_n XOR B_n and A_n AND B_n suffices to compute S_n (the sum) and C_n (the carry out) given C_{n-1} (carry input from the previous stage). Note that the sensing schemes discussed enable us to perform the bitwise XOR and AND operations simultaneously, thereby performing an ADD operation with a single array access. Figure 5 also expresses the ADD operation in terms of the outputs of bitwise operations, O_{AND} and O_{XOR}. Three additional logic gates are required to enable this computation.

Full-adder function	$S_n = (A_n \text{ XOR } B_n) \text{ XOR } C_{n-1}$ (bitwise operation) $C_n = ((A_n \text{ XOR } B_n) \text{ AND } C_{n-1}) \text{ OR } (A_n \text{ AND } B_n)$
In-Memory ADD	$S_n = O_{XOR} \text{ XOR } C_{n-1}$ $C_n = (O_{XOR} \text{ AND } C_{n-1}) \text{ OR } O_{AND}$

Fig. 5: In-Memory ADD operation

B. STT-CiM array

In this section, we present the array-level design of STT-CiM using the mechanisms described above. As shown in Figure 6, the proposed STT-CiM memory array takes an additional input CiMType that indicates the type of compute-in-memory operation that needs to be performed for every memory access. The CiM decoder interprets this input and generates appropriate control signals to perform the desired logic operation. In order to enable compute-in-memory operations, the read peripheral circuits present in each column (sensing circuit and global reference generation circuit in Figure 6) are enhanced, while the core data array remains the same as in standard STT-MRAM. The address decoder is enhanced so as to enable multiple wordlines for CiM operations. The write peripheral circuits remain unchanged, as write operations are identical to standard STT-MRAM. We next describe the enhancements to the read peripheral circuits.

Sensing circuitry. Figure 6 shows the sensing circuit enhanced to support all the logic operations discussed in Section IV-A. It consists of two sense amplifiers, a CMOS NOR gate, three multiplexers and three additional logic gates for ADD operation. We note that the area and power overheads associated with these enhancements are minimal since the sensing circuit constitutes a small fraction of the total memory area/power. As shown in the figure, the reference currents (I_{refl}, I_{refr}) produced by the global reference generation circuit are fed to the two sense amplifiers in order to realize the sensing schemes discussed in Section IV-A. The three MUX control signals (sel_0, sel_1, sel_2) are generated by the CiM decoder to select the desired compute-in-memory operation.

Reference generation. Figure 6 illustrates the modified reference generation circuit used to produce the additional reference currents necessary for the proposed sensing schemes. It includes two reference stacks, one each for the two sense amplifiers in the sensing circuit. Each stack consists of three bit-cells programmed to offer resistances R_P, R_{AP} and R_{REF}, respectively. R_{REF} [2] represents the read reference MTJ used in a standard STT-MRAM memory array. The CiM decoder generates control signals (rwl_0, rwl_1, ... rwr_1, rwr_2) that are used to enable a subset of these bit-cells in the reference stacks, which in turn produces the desired reference currents.

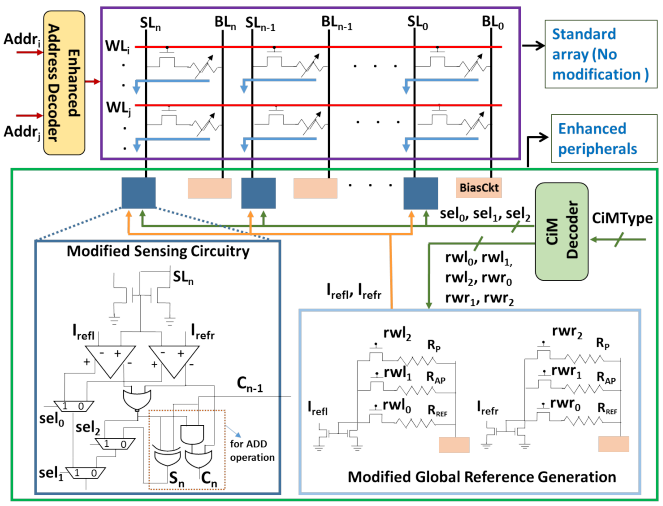

Fig. 6: STT-CiM array structure

The STT-CiM array can perform both regular memory operations and a range of CiM operations. The normal read operation is performed by enabling a single wordline and setting sel_0, sel_1, and rwl_0 to logic '1'. On the other hand, a CiM operation is performed by enabling two wordlines and setting CiMType to the appropriate value, which results in computing the desired function of the enabled words.

C. CiM operation under process-variations

The STT-CiM array suffers from the same failure mechanisms, *viz.* read disturb failures, read decision failures and write failures, that are observed in a standard STT-MRAM memory array. In this section, we compare the failure rates in STT-CiM and standard STT-MRAM. Normal read/write operations in STT-CiM have the same failure rate as in a standard STT-MRAM, since the read/write mechanism is identical. However, CiM operations differ in their failure rates, since the currents that flow through each bit-cell differ when enabling two wordlines simultaneously. In order to analyze the read disturb and read decision failures under process variations for CiM operations, we performed a Monte-carlo circuit-level simulation on 1 million samples considering variations in MTJ oxide thickness (σ/μ = 2 %), transistor V_T (σ/μ = 5%), and MTJ cross sectional area (σ/μ = 5%). Figure 7 shows the probability density distribution of the possible currents obtained during read and CiM operations on these 1 million samples.

CiM disturb failures. As shown in Figure 7, the overall current flowing through the source line is slightly higher in case of a CiM operation as compared to a normal read. However this increased current is divided between the two parallel paths, and consequently the net read current flowing through each bit-cell (MTJ) is reduced. Hence, the read disturb failure rate is even lower for CiM operations than normal read operations.

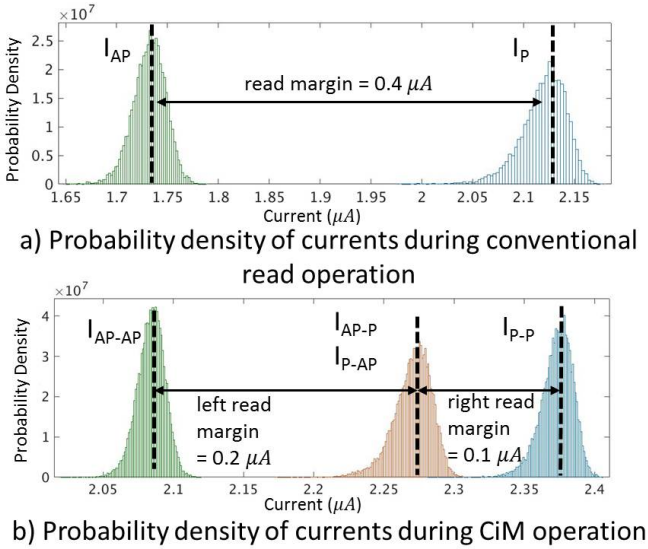

Fig. 7: Probability density distribution of I_{SL} under process variations during read and CiM operations

CiM decision failures. The net current flowing through the source line (I_{SL}) in case of a CiM operation can have 3 possible values, *i.e.*, I_{P-P}, I_{AP-P} (I_{P-AP}), I_{AP-AP}. A read decision failure occurs during a CiM operation when the current I_{P-P} is interpreted as I_{AP-P} (or vice versa), or when I_{AP-AP} is inferred as I_{AP-P} (or vice versa). In contrast to normal reads, CiM operations have two read margins (shown in Figure 7(b)), one between I_{P-P} and I_{AP-P} and another between I_{AP-P} and I_{AP-AP}. Our simulation results show that the read margins for CiM operations are lower as compared to normal reads, therefore they are more prone to decision failures. Moreover, the read margins in CiM operations are unequal [3], thus we have more failures arising due to the read margin between I_{P-P} and I_{AP-P}.

ECC for STT-CiM. In order to mitigate these failures in STT-MRAM, various ECC schemes have been previously explored [13], [14]. We show that ECC techniques that provide single error correction and double error detection (SECDED) and double error correction and triple error detection (DECTED) can be used to address the decision failures in CiM operations as well. This is feasible because the codeword properties for these codes are retained for a CiM XOR operation. Figure 8 shows the codeword retention property of a CiM XOR operation using a simple Hamming code. As shown

[2] $R_{AP} > R_{REF} > R_P$

[3] Although resistances may be equally separated, the currents are not since they depend inversely on resistance.

in the figure, $word_1$ and $word_2$ are augmented with ECC bits (p_1, p_2, p_3) and stored in memory as $InMemW_1$ and $InMemW_2$ respectively. A CiM XOR operation performed on these stored words ($InMemW_1$, $InMemW_2$) results in the ECC codeword for $word_1$ XOR $word_2$, therefore the codewords are preserved for CiM XORs. We leverage this retention property of CiM XORs to detect and correct errors in *all CiM operations*. This is enabled by the fact that our design STT-CiM always computes bitwise XOR (CiM XOR) irrespective of the desired CiM operation. We demonstrate the proposed error detection and correction mechanism for CiM operations in Figure 9. Let us assume that data bit d_1 suffers from decision failure during CiM operations, as shown in the figure. As a result, logic '11' (I_{P-P}) is inferred as logic '10' (I_{AP-P}) which leads to erroneous CiM outputs. An error detection logic operating on the CiM XOR output (shown in Figure 9) detects an error in the d_1 data bit which can be corrected directly for a CiM XOR operation by simply flipping the erroneous bit. However for other CiM operations that do not retain codewords, we perform two sequential normal reads on words $InMemW_1$ and $InMemW_2$, and correct the erroneous bits by recomputing them using an error detection and correction unit (discussed in section V). Note that, such corrections lead to overheads, as we need to access memory array 3 times (compared to 2 times in STT-MRAM), thereby reducing the efficiency of STT-CiM. However, our variation analysis shows that corrections on CiM operations are infrequent, leading to overall improvements.

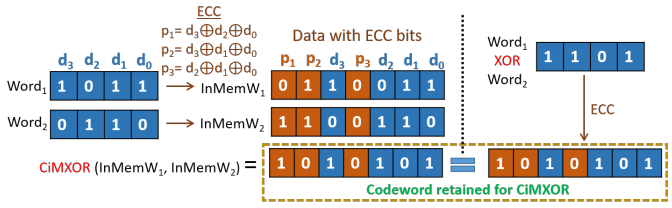

Fig. 8: Codeword retention property of CIM XOR

ECC design methodology. We use the methodology employed in [13] to determine ECC requirements for both STT-MRAM (baseline) as well as the proposed STT-CiM design. The approach uses circuit level simulations to determine the bit-level error probability, which is then used to estimate the array level yield. Moreover, the ECC scheme is selected based on the target yield requirement. Our simulation shows that 1 bit failure probability of normal reads and CiM operations are 4.2x10^{-8} and 6x10^{-5} respectively. With these obtained bit-level failure rates and assuming a target yield of 99%, the ECC requirement for 1MB STT-MRAM is single error correction and double error detection (SECDED), whereas for 1MB STT-CiM is three error correction and four error detection (3EC4ED). Note that the overheads of the ECC schemes are fully considered and reflected in our experimental results. Moreover, our simulation shows that the probability of CiM operations having errors is 0.1, *i.e.*, 1 in 10 CiM operations will have an error, which will always be detected by using 3EC4ED code on CiM XORs. Further, these errors are directly corrected for CiM XORs using 3EC4ED code, and by reverting to near-memory computation for other CiM operations.

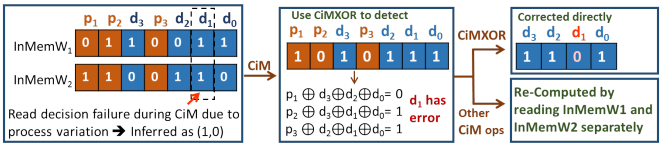

Fig. 9: Error detection and correction for CiM operations

V. STT-CiM Architecture

In order to evaluate the application-level benefits of STT-CiM, we integrate it as a scratchpad in the memory hierarchy of a programmable processor [52]. This section describes optimizations to increase the efficiency of STT-CiM and architectural enhancements required to expose it to software.

A. Optimizations for STT-CiM

In order to further the efficiency improvements obtained by STT-CiM, we propose additional optimizations.

Vector CiM operations. Many modern computing workloads exhibit significant data parallelism. To further enhance the efficiency of STT-CiM for such applications, we introduce Vector Compute-in-Memory (VCiM) operations. The key idea behind VCiM operations is to exploit the internal memory bandwidth to perform CiM operations on the elements of a vector simultaneously. Figure 10 shows how the internal memory bandwidth (32xN bits) can be significantly larger than the limited I/O bandwidth (32 bits) visible to the processor. We exploit this internal bandwidth to perform vector operations (N words wide) within STT-CiM. However, the computed vector cannot be directly transferred to processor in one access due to I/O limitations.

To address this issue, we leverage the reduction computation patterns commonly observed in several applications and introduce a Reduce Unit (RU) shown in Figure 10. The RU takes an array of data elements as inputs and reduces it to a single data element. A reduce unit can support various reduction operations such as summation, zero-compare, etc. (shown in Table II). The overheads of the RU depend on two factors: (i) the number of different reduction operations supported, and (ii) the maximum vector length allowed (can be between 2 to N words). In our evaluation, we have included the timing and area overheads associated with the proposed RU. To limit the area overheads, we support only two common reduction operations, *i.e.*, summation and zero-compare, and evaluate our design for vector lengths of 4 and 8. Consider the computation of $\sum_{i=1}^{N} A[i] + B[i]$, where arrays A and B are stored in rows i and j respectively (shown in Figure 10). To compute the desired function using a VCiM operation, we activate rows i and j simultaneously, and configure the sensing circuitry to perform an ADD operation and the RU to perform accumulation of the resulting output. Note that the summation would require 2N memory accesses in a conventional memory. With scalar CiM operations, it would require N memory

accesses. With the proposed VCiM operations, only a single memory access is required.

In order to realize a wider variety of in-memory operations, we further enhance the proposed RU with a low overhead compute unit. Specifically, this unit enables operations such as Euclidean distance and L1/L2 norms to be computed by STT-CiM.

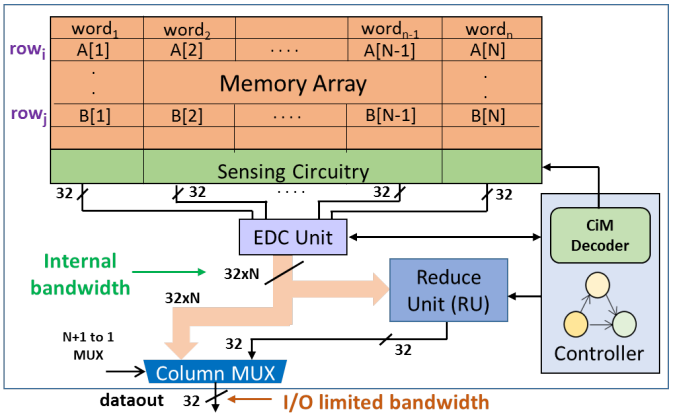

Fig. 10: STT-CiM supporting In-memory vector operation

Error Detection and Correction. To enable correction of erroneous bits for CiM operations, we introduce an Error Detection and Correction (EDC) unit which implements the 3EC4ED ECC scheme. The EDC unit checks for errors using the CiM XOR output and signals the controller (shown in Figure 10) upon detection of erroneous computations. The controller on receiving this error detection signal performs the required corrective actions.

TABLE II: Examples of reduction operations

Type	Function RuOut = f (IN_1, IN_2,...,IN_N)
Summation	RuOut=IN_1 + IN_2+ .. + IN_N
Zero-Compare	RuOut[k] = (IN_k == 0) ? 0 : 1

B. Architectural Extensions for STT-CiM

To integrate STT-CiM in a programmable processor based system, we propose the following architectural enhancements.

ISA extension. We propose extensions to the ISA of a programmable processor to support CiM operations. To this end, we introduce a set of new instructions in the ISA (CiMXOR, CiMNOT,CiMAND, CiMADD ...) that are used to invoke the different types of operations that can be performed in the STT-CiM array. In a load instruction, the requested address is sent to the memory, and the memory returns the data stored at the addressed location. However, in the case of a CiM instruction, the processor is required to provide addresses of two memory locations instead of a single one, and the memory operates on the two data values to return the final output.

Format: **Opcode** Reg1 Reg2 Reg3
Example: **CiMXOR** R_{ADDR1} R_{ADDR2} R_{DEST} (1)

Equation 1 shows the format of a CiM instruction with an example. As shown, both the addresses required to perform CiMXOR operations are provided through registers. The format is similar to a regular arithmetic instruction that accesses two register values, performs the computation, and stores the result back in a register.

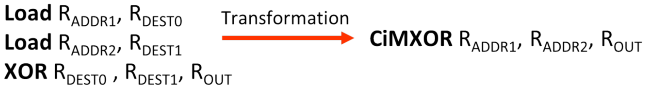

Fig. 11: Program transformation for CiMXOR

Program transformation. To exploit the proposed CiM instructions at the application-level, an assembly-level program transformation is performed, wherein specific sequences of instructions in the compiled program are mapped to suitable CiM instructions in the ISA. Figure 11 shows an example transformation where two load instructions followed by an XOR instruction can be mapped to a single CiMXOR instruction.

Bus and interface support. In a programmable processor based system, the processor and the memory communicate via a system bus or on-chip network [4]. This makes it essential to analyse the impact of CiM operations on the bus and the corresponding bus interface. As discussed above, a CiM operation is similar to a load instruction with the key difference that it sends two addresses to the memory. Conventional system buses only allow sending a single address onto the bus via the address channel. In order to send the second address for CiM operations, we utilize the unused writedata channel of the system bus, which is unutilized during a CiM operation. Besides the two addresses, the processor also sends the type of CiM operation (CIMType) that needs to be performed. Note that it may be possible to overlay the CIMType signal onto the existing bus control signals; however, such optimizations strongly depend on the specifics of the bus protocol being used. In our design, we assume that 3 control bits are added to the bus and account for the resulting overheads in our experiments.

C. Data Mapping

In order to perform a CiM instruction, the locations of its operands in memory must satisfy certain constraints. Let us consider a memory organization consisting of several banks where each bank is an array that contains rows and columns. In this case, a CiM operation can be performed on two data elements only if they satisfy three key criteria: (i) they are stored in the same bank, (ii) they are mapped to different rows, and (iii) they are aligned and placed in the same set of columns.

Consequently, a suitable data placement technique is required that maximizes the benefits of STT-CiM. We observe that the target applications for STT-CiM have well defined computation patterns, facilitating such a data placement. Figure 12 shows three general computation patterns observed in these target applications. We next discuss these compute patterns and the corresponding data placement techniques.

[4]While we consider the case of a shared bus for illustration, the same enhancements can be applied to more complex networks.

Type I. This pattern, shown in the top row of Figure 12, involves element-to-element operations (OPs) between two arrays, *e.g.*, A and B. In order to effectively utilize STT-CiM for this compute pattern, we utilize the *array alignment* technique (shown in Figure 12(a)) that ensures alignment of elements A[i] and B[i] of arrays A and B for any value of i. This enables the conversion of operation A[i] OP B[i] into a CiM operation. An extension to this technique is the *row-interleaved placement* shown in Figure 12(b). This technique is applicable to larger data structures that do not fully reside in same memory bank. It ensures that both the elements, *i.e.*, A[i] and B[i], are mapped to the same bank for any value of i, and satisfy the alignment criteria for a CiM operation.

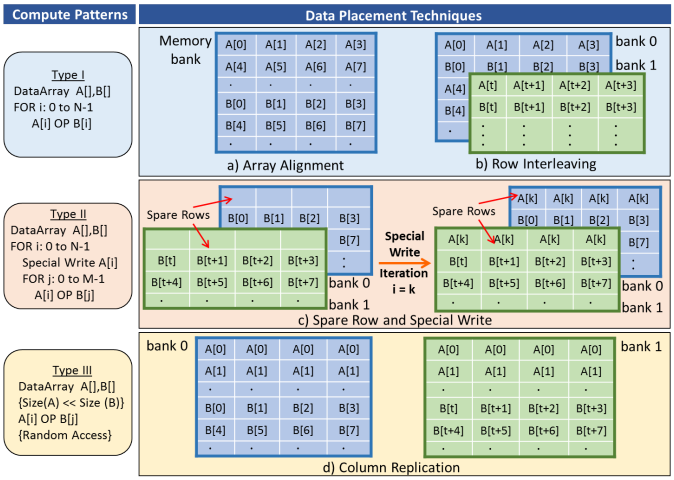

Fig. 12: Data mapping for various computation patterns

Type II. This pattern, shown in the middle row of Figure 12, involves a nested loop in which the inner loop iteration consists of a single element of array A being operated with several elements of array B. For this one-to-many compute pattern, we introduce a *spare row* technique for STT-CiM data alignment. In this technique, a spare row is reserved in each memory bank to store copies of an element of A. As shown in Figure 12(c), in the k^{th} iteration of the outer for-loop, a special write operation is used to fill the spare rows in all banks with A[k]. This results in each element of array B becoming aligned with a copy of A[k], thereby allowing CiM operations to be performed on them. Note that the special write operation introduces energy and performance overheads, but this overhead is amortized over all inner loop iterations, and is observed to be quite insignificant in our evaluations.

Type III. In this pattern, shown in the bottom row of Figure 12, operations are performed on an element drawn from a small array A and an element from a much larger array B. The elements are selected arbitrarily, *i.e.*, without any predictable pattern. For example, consider when a small sequence of characters needs to be searched within a much larger input string. For this pattern, we propose a *column replication* technique to enable CiM operations, as shown in Figure 12(d). In this technique, a single element of the small array A is replicated across columns to fill an entire row. This ensures that each element of A is aligned with every element of B, enabling a CiM operation to be utilized. Note that the initial overhead due to data replication is very small, as it pales in comparison to the number of memory accesses to the larger array.

VI. Experimental Methodology

In this section, we discuss the device-to-architecture simulation framework (shown in Figure 13) used in our experiments to evaluate the performance and energy benefits of STT-CiM at the array-level and system-level.

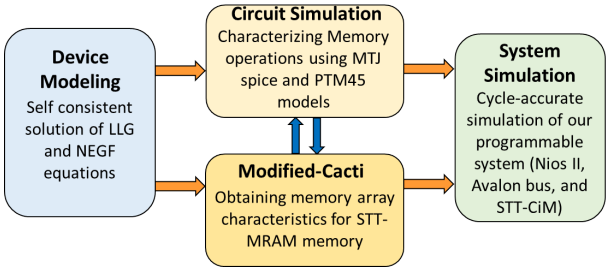
Device Modeling
Self consistent solution of LLG and NEGF equations
Circuit Simulation
Characterizing Memory operations using MTJ spice and PTM45 models
Modified-Cacti
Obtaining memory array characteristics for STT-MRAM memory
System Simulation
Cycle-accurate simulation of our programmable system (Nios II, Avalon bus, and STT-CiM)

Fig. 13: STT-CiM device-to-architecture simulation framework

Device/Circuit modeling. We first characterize the bit-cells using SPICE-compatible MTJ models that are based on self-consistent solution of Landau-Lifshitz-Gilbert (LLG) magnetization dynamics and Non-Equilibrium-Green's Function (NEGF) electron transport [55]. Table III shows the MTJ device parameters [56] used in our experiments. Using 45nm bulk CMOS technology and the MTJ models, the memory array along with the associated peripherals and extracted parasitics was simulated in SPICE for various operations (read, write and CiM operations) to obtain array-level timing and energy characteristics. The obtained characteristics were then used as technology parameters in a modified version of CACTI [57] that is capable of estimating system-level properties for a spin-based memory. The variation analysis to compute failure rates was performed considering variations in MTJ oxide thickness (σ/μ = 2 %), transistor V_T (σ/μ = 5%), and MTJ cross sectional area (σ/μ = 5%).

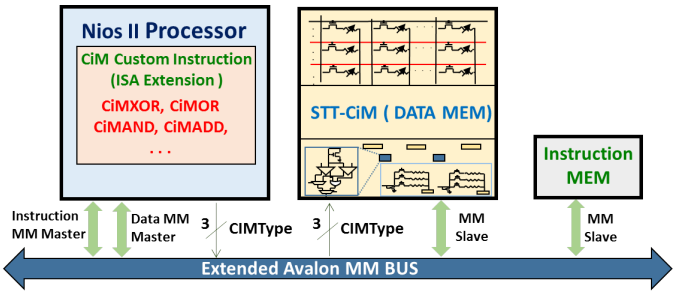

Fig. 14: System level integration of STT-CiM

System level simulation. We evaluated STT-CiM as a 1MB scratchpad for an Intel Nios II processor [52]. Figure 14 shows the integration of STT-CiM in the memory hierarchy of the programmable processor. In order to expose the STT-CiM operations to software, we extended the Nios II processor's instruction set with custom instructions. The Avalon on-chip

bus was also extended to support CiM operations. Cycle-accurate RTL simulation was used to obtain the execution time and the memory access traces for various benchmarks. These traces along with the energy results obtained through the modified CACTI tool were used to estimate the total memory energy.

TABLE III: Device parameters

Material System	Ta/CoFeB/MgO
MTJ Type	PMA
Saturation Magnetization(M_S)	1.58T
Damping Factor, (α)	0.028
Polarization	0.62
Interface Anisotropy	1.3mJ/m^2
MTJ Dimension	40nm x 40nm x 1.32nm
Oxide Thickness (t_{ox})	1.1nm
Energy Barrier	65KT
T	300K
RA Product	18ohm-μm^2
TMR	124%
CMOS Technology	45nm Bulk CMOS
Assumed Variation (σ/μ)	t_{ox} = 2%, MTJ Area = 5% transistor V_T=5%

TABLE IV: Benchmark applications

Applications	Domain
Aho-Corasick (AHC)	String matching
Knuth-Morris-Pratt (KMP)	String matching
Edit Distance (EDIST)	Text processing, Computational biology
Bit-Blit (BLIT)	Low-level graphics
Longest Common Subsequence (LCS)	Data compression, Bio-informatics
Rivest Cipher 4 (RC4)	Cryptography
IMGSEG	Image segmentation
GLVQ	Eye detection
K-means clustering (KMEANS)	Digit recognition
Optical Character Recognition (OCR)	Character recognition
Multi Layer Perceptron (MLP)	Protein structure classification
Support Vector Machines (SVM)	Text classification

Benchmark applications. We evaluate STT-CiM on a suite of twelve applications drawn from various domains (Table IV).

VII. Results

In this section, we first present an array-level analysis of STT-CiM and then quantify its benefits through system-level energy and performance evaluation.

A. Array-level analysis

Energy. Figure 15 shows the energy consumed by standard read and a representative CiM operation (XOR) in a 1MB STT-CiM array (second and third bars). Each bar shows the energy breakdown into the major components, *i.e.*, peripheral circuitry (PeriphCkt), wordline (WordL), bitline (BitL), reference generation circuitry (REF), sense amplifier (SenseA), and error correction circuitry (ECC). For reference, we provide the read energy for an STT-MRAM array of the same capacity (first bar) and all energy numbers are normalized to this value. A normal read operation in STT-CiM incurs an energy overhead of about 4.4%, which arises primarily due to the extra peripheral circuits (PeriphCkt) and more stringent ECC requirements. Our STT-CiM design uses a 3EC4ED ECC scheme (as compared to SECDED in STT-MRAM) that accounts for about 3% of the 4.4% energy overhead. The XOR CiM operation consumes higher energy than a standard read operation mainly due to the charging of multiple wordlines and a slightly higher source line current. Since a CiM operation substitutes two normal read operations, we also present the energy required for two reads in a standard STT-MRAM (last bar). Note that an array-level comparison greatly understates the benefits of STT-CiM [5]. Nevertheless, it is worth noting that even at the array level, STT-CiM consumes 38.2% less energy. The benefits mainly arise from a lower bitline dynamic energy (BitL), since only a single access to the memory array is required for STT-CiM.

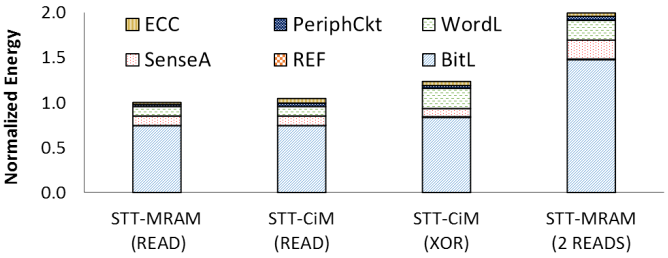

Fig. 15: Array-level energy evaluation of STT-CiM

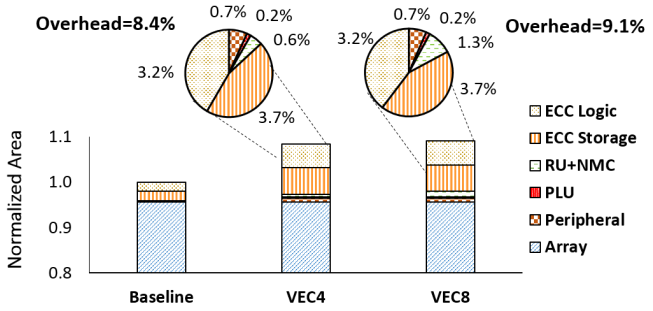

Fig. 16: Array-level area evaluation of STT-CiM

Area and access time. Figure 16 shows the area breakdown for two STT-CiM designs that support vector operations of length 4 and 8 (VEC4 and VEC8 respectively). As compared to our baseline the area overheads for VEC4 and VEC8, are 8.4% and 9.1%, respectively. As shown in Figure 16, ECC storage and ECC Logic forms the major component of the area overheads (3.7% and 3.2% respectively). Note that, for STT-CiM designs the total area is dominated by the core array which remains unchanged. Finally, the access time overhead for STT-CiM was found to be only ∼0.8%, because the wordline and bitline delays dominate the total memory latency.

[5]It does not consider the major impact of reduced data transfers between the processor and memory, which is considered in the system-level results presented in the next subsection.

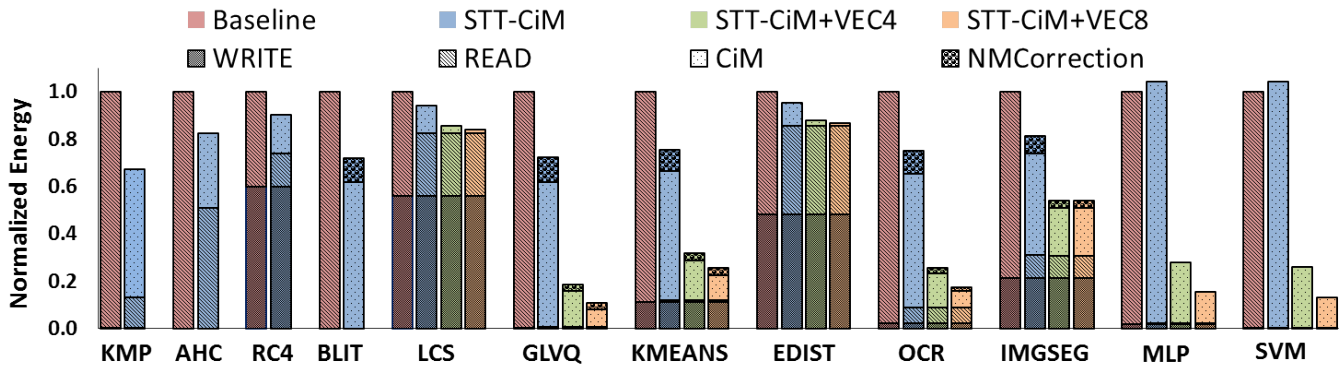

Fig. 17: Application-level memory energy

B. Application-level memory energy

We next present the system level memory energy benefits of using STT-CiM in the programmable processor based system described in Figure 14. We evaluated the total memory energy consumed by STT-CiM across the application benchmarks, and compared it with our baseline design that uses standard STT-MRAM. Figure 17 shows the breakdown of different energy components, *viz.* Read, Write and CiM, that contribute to the overall memory energy in both the STT-MRAM and proposed STT-CiM designs. In addition, it also shows the energy overheads due to near memory corrections (NMCorrections) on failing CiM operations. The total memory energy for an application is normalized to the memory energy consumed by the baseline design. For the proposed STT-CiM design, we evaluated a version without vector operations (STT-CiM), and two versions with vector length 4 and 8 (STT-CiM+VEC4 and STT-CiM+VEC8, respectively). Across all benchmarks, we observe 1.26x, 2.77x and 3.83x average improvement in energy for STT-CiM, STT-CiM+VEC4 and STT-CiM+VEC8, respectively.

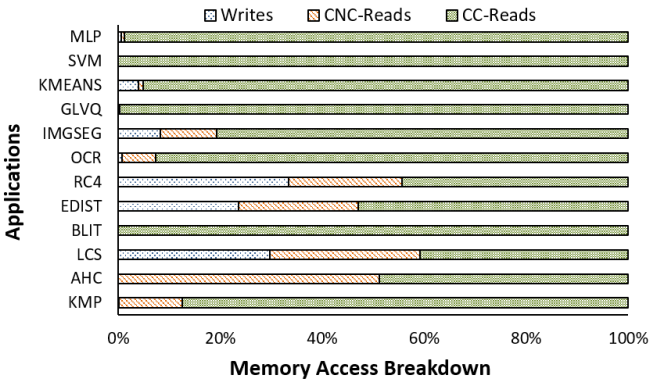

Fig. 18: Memory access breakdown

To provide further insights into the energy benefits, Figure 18 presents a breakdown for memory accesses made by each application into 3 categories — writes, reads that cannot be converted into CiM operations (CiM non-convertible reads or CNC-Reads), and CiM convertible reads (CC-Reads). We see that applications where CC-Reads dominate the total memory accesses (KMP, BLIT, GLVQ, KMEANS, OCR, IMGSEG, MLP, SVM in Figure 18) exhibit higher energy benefits from STT-CiM (Figure 17). Among these applications, those that benefit from vectorization achieve the highest savings (GLVQ, KMEANS, OCR, MLP, SVM). Applications with relatively fewer CC-Reads or more frequent writes (AHC, LCS, RC4, EDIST) exhibit relatively lower energy savings. CNC-Reads and writes are not benefited by STT-CiM, and writes in particular consume significantly (∼3x) higher energy than reads.

C. System-level performance

Figure 19 shows the speedup for the Nios II processor system integrated with STT-CiM across various applications. The speedup shown in the figure is with respect to the baseline design, *i.e.*, the processor system integrated with a standard STT-MRAM based memory. As discussed in Section V-B, CiM lowers the total number of memory accesses as well as the number of instructions executed, which leads to performance benefits at the system level. Overall, for STT-CiM without vector operation, we observe performance benefits ranging from 1.07X to 1.36X. With vector operations, the average speedup increased to 3.25x and 3.93x for vector lengths of 4 and 8, respectively. Comparing Figures 18 and 19, we see that the factors that indicate higher energy savings for an application (large fraction of memory accesses are CC-Reads, opportunities for vectorization exist) are also predictive of higher performance improvements.

In order to demonstrate the performance sensitivity to memory latency, we vary the memory latency and evaluate the execution time for each application. Figure 20 shows the results of this sensitivity analysis. On the Y-axis, we have the speedup, and on the X-axis the memory latency. We observe that increasing memory latency yields higher performance benefits for STT-CiM. This is expected, because the CiM operations lower the number of total memory accesses, as a result, a higher memory latency will result in a larger performance benefit. On an average, we achieve 1.13x speedup for a memory latency of 1 cycle, and 1.26x speedup for a memory latency of 16 cycles, thereby illustrating the effectiveness of the proposed approach.

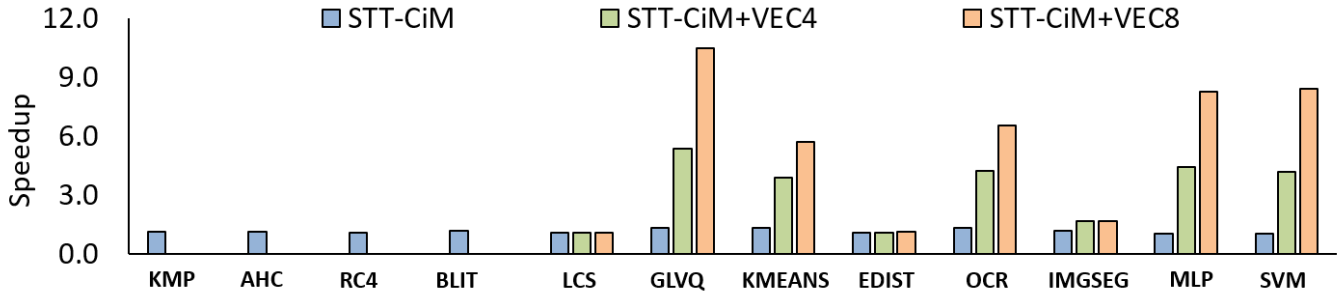

Fig. 19: Application-level system performance

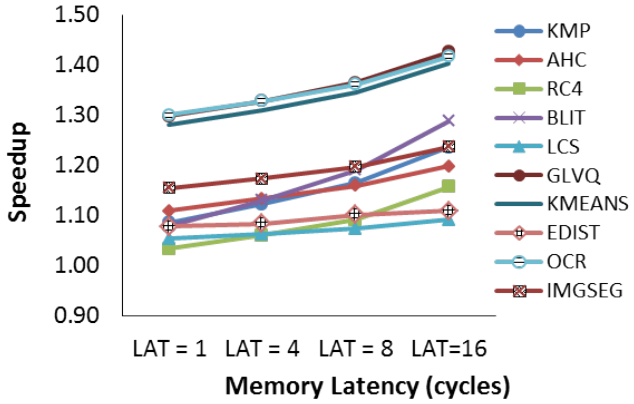

Fig. 20: Performance sensitivity to memory latency

VIII. Conclusion

STT-MRAM is a promising candidate for future on-chip memories. In this work, we proposed STT-CiM, an enhanced STT-MRAM that can perform a range of arithmetic, logic and vector compute-in-memory operations. We addressed a key challenge associated with these in-memory operations, *i.e.* reliable computation under process variations. We utilized the proposed design (STT-CiM) as a scratchpad in the memory hierarchy of a programmable processor, and introduced ISA extensions and on-chip bus enhancements to support in-memory computations. We proposed architectural optimizations and data mapping techniques to enhance the efficiency of STT-CiM. A device-to-architecture simulation framework was used to evaluate the benefits of STT-CiM. Our experiments indicate that STT-CiM achieves substantial improvements in energy and performance, and shows considerable promise in alleviating the processor-memory gap.

References

[1] http://www.everspin.com/ .

[2] Dmytro Apalkov, Alexey Khvalkovskiy, Steven Watts, Vladimir Nikitin, Xueti Tang, Daniel Lottis, Kiseok Moon, Xiao Luo, Eugene Chen, Adrian Ong, Alexander Driskill-Smith, and Mohamad Krounbi. Spin-transfer Torque Magnetic Random Access Memory (STT-MRAM). *J. Emerg. Technol. Comput. Syst.*, 9(2):13:1–13:35, May 2013.

[3] http://www.avalanche-technology.com/ .

[4] A. Jog, A. K. Mishra, C. Xu, Y. Xie, V. Narayanan, R. Iyer, and C. R. Das. Cache revive: Architecting volatile STT-RAM caches for enhanced performance in CMPs. In *Design Automation Conference (DAC), 2012 49th ACM/EDAC/IEEE*, pages 243–252, June 2012.

[5] Ping Zhou, Bo Zhao, Jun Yang, and Youtao Zhang. Energy Reduction for STT-RAM Using Early Write Termination. In *Proceedings of the 2009 International Conference on Computer-Aided Design*, ICCAD '09, pages 264–268, November 2009.

[6] S. Chatterjee, M. Rasquinha, S. Yalamanchili, and S. Mukhopadhyay. A Scalable Design Methodology for Energy Minimization of STTRAM: A Circuit and Architecture Perspective. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 19(5):809–817, May 2011.

[7] Yusung Kim, Sumeet Kumar Gupta, Sang Phill Park, Georgios Panagopoulos, and Kaushik Roy. Write-optimized Reliable Design of STT MRAM. In *Proceedings of the 2012 ACM/IEEE International Symposium on Low Power Electronics and Design*, ISLPED '12, pages 3–8, New York, NY, USA, 2012. ACM.

[8] H. Noguchi, K. Ikegami, K. Kushida, K. Abe, S. Itai, S. Takaya, N. Shimomura, J. Ito, A. Kawasumi, H. Hara, and S. Fujita. 7.5 A 3.3ns-access-time 71.2 uW/MHz 1Mb embedded STT-MRAM using physically eliminated read-disturb scheme and normally-off memory architecture. In *2015 IEEE International Solid-State Circuits Conference - (ISSCC) Digest of Technical Papers*, pages 1–3, Feb 2015.

[9] Sang Phill Park, Sumeet Gupta, Niladri Mojumder, Anand Raghunathan, and Kaushik Roy. Future cache design using STT MRAMs for improved energy efficiency: Devices, circuits and architecture. In *Proceedings of the Design Automation Conference*, pages 492–497, June 2012.

[10] C. W. Smullen, V. Mohan, A. Nigam, S. Gurumurthi, and M. R. Stan. Relaxing non-volatility for fast and energy-efficient STT-RAM caches. In *Proceedings of the International Symposium on High Performance Computer Architecture*, pages 50 –61, February 2011.

[11] W. Xu, H. Sun, X. Wang, Y. Chen, and T. Zhang. Design of Last-Level On-Chip Cache Using Spin-Torque Transfer RAM (STT RAM). *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 19(3):483–493, March 2011.

[12] A. Aziz, N. Shukla, S. Datta, and S. K. Gupta. COAST: Correlated material assisted STT MRAMs for optimized read operation. In *2015 IEEE/ACM International Symposium on Low Power Electronics and Design (ISLPED)*, pages 1–6, July 2015.

[13] K. W. Kwon, X. Fong, P. Wijesinghe, P. Panda, and K. Roy. High-Density and Robust STT-MRAM Array Through Device/Circuit/Architecture Interactions. *IEEE Transactions on Nanotechnology*, 14(6):1024–1034, Nov 2015.

[14] B. Del Bel, J. Kim, C. H. Kim, and S. S. Sapatnekar. Improving STT-MRAM density through multibit error correction. In *2014 Design, Automation Test in Europe Conference Exhibition (DATE)*, pages 1–6, March 2014.

[15] Wang Kang, Liuyang Zhang, Weisheng Zhao, J.-O. Klein, Youguang Zhang, D. Ravelosona, and C. Chappert. Yield and Reliability Improvement Techniques for Emerging Nonvolatile STT-MRAM. *Emerging and Selected Topics in Circuits and Systems, IEEE Journal on*, 5(1):28–39, March 2015.

[16] Xuanyao Fong, Yusung Kim, S.H. Choday, and K. Roy. Failure Mitigation Techniques for 1T-1MTJ Spin-Transfer Torque MRAM Bitcells. *IEEE Trans. VLSI Systems*, 22(2):384–395, Feb 2014.

[17] G. S. Kar, W. Kim, T. Tahmasebi, J. Swerts, S. Mertens, N. Heylen, and T. Min. Co/Ni based p-MTJ stack for sub-20nm high density stand alone and high performance embedded memory application. In *2014 IEEE International Electron Devices Meeting*, pages 19.1.1–19.1.4, Dec 2014.

[18] Ashish Ranjan, Swagath Venkataramani, Xuanyao Fong, Kaushik Roy, and Anand Raghunathan. Approximate Storage for Energy Efficient Spintronic Memories. In *Proceedings of the 52Nd Annual Design*

Automation Conference, DAC '15, pages 195:1–195:6, New York, NY, USA, 2015. ACM.

[19] A. K. Mishra, X. Dong, G. Sun, Y. Xie, N. Vijaykrishnan, and C. R. Das. Architecting on-chip interconnects for stacked 3D STT-RAM caches in CMPs. In *2011 38th Annual International Symposium on Computer Architecture (ISCA)*, pages 69–80, June 2011.

[20] K. Lee and S. H. Kang. Development of Embedded STT-MRAM for Mobile System-on-Chips. *IEEE Transactions on Magnetics*, 47(1):131–136, Jan 2011.

[21] A. Nigam, C. W. Smullen, V. Mohan, E. Chen, S. Gurumurthi, and M. R. Stan. Delivering on the promise of universal memory for spin-transfer torque RAM (STT-RAM). In *IEEE/ACM International Symposium on Low Power Electronics and Design*, pages 121–126, Aug 2011.

[22] A. Jadidi, M. Arjomand, and H. Sarbazi-Azad. High-endurance and performance-efficient design of hybrid cache architectures through adaptive line replacement. In *IEEE/ACM International Symposium on Low Power Electronics and Design*, pages 79–84, Aug 2011.

[23] Y. Zhang, W. Zhao, J. O. Klein, W. Kang, D. Querlioz, C. Chappert, and D. Ravelosona. Multi-level cell Spin Transfer Torque MRAM based on stochastic switching. In *2013 13th IEEE International Conference on Nanotechnology (IEEE-NANO 2013)*, pages 233–236, Aug 2013.

[24] J. Zhao and Y. Xie. Optimizing bandwidth and power of graphics memory with hybrid memory technologies and adaptive data migration. In *2012 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, pages 81–87, Nov 2012.

[25] Wei Xu, Yiran Chen, Xiaobin Wang, and Tong Zhang. Improving STT MRAM Storage Density Through Smaller-than-worst-case Transistor Sizing. In *Proceedings of the 46th Annual Design Automation Conference*, DAC '09, pages 87–90, New York, NY, USA, 2009. ACM.

[26] M. Rasquinha, D. Choudhary, S. Chatterjee, S. Mukhopadhyay, and S. Yalamanchili. An energy efficient cache design using Spin Torque Transfer (STT) RAM. In *2010 ACM/IEEE International Symposium on Low-Power Electronics and Design (ISLPED)*, pages 389–394, Aug 2010.

[27] Kostas Pagiamtzis and Ali Sheikholeslami. Content-addressable memory (CAM) circuits and architectures: A tutorial and survey. *IEEE Journal of Solid-State Circuits*, 41(3):712–727, March 2006.

[28] M. Kang, M. S. Keel, N. R. Shanbhag, S. Eilert, and K. Curewitz. An energy-efficient VLSI architecture for pattern recognition via deep embedding of computation in SRAM. In *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8326–8330, May 2014.

[29] J.P. Wang and J.D. Harms. General structure for computational random access memory (CRAM), November 13 2014. US Patent App. 14/259,568.

[30] Shuangchen Li, Cong Xu, Qiaosha Zou, Jishen Zhao, Yu Lu, and Yuan Xie. Pinatubo: A processing-in-memory architecture for bulk bitwise operations in emerging non-volatile memories. In *Proceedings of the 53rd Annual Design Automation Conference*, DAC '16, pages 173:1–173:6, New York, NY, USA, 2016. ACM.

[31] V. Seshadri, K. Hsieh, A. Boroum, D. Lee, M. A. Kozuch, O. Mutlu, P. B. Gibbons, and T. C. Mowry. Fast Bulk Bitwise AND and OR in DRAM. *IEEE Computer Architecture Letters*, 14(2):127–131, July 2015.

[32] Qing Guo, Xiaochen Guo, Ravi Patel, Engin Ipek, and Eby G. Friedman. AC-DIMM: Associative Computing with STT-MRAM. In *Proceedings of the 40th Annual International Symposium on Computer Architecture*, ISCA '13, pages 189–200, New York, NY, USA, 2013. ACM.

[33] Jintao Zhang, Zhuo Wang, and Naveen Verma. A machine-learning classifier implemented in a standard 6T SRAM array. In *VLSI Circuits (VLSI-Circuits), 2016 IEEE Symposium on*, pages 1–2. IEEE, June 2016.

[34] B. Falsafi, M. Stan, K. Skadron, N. Jayasena, Y. Chen, J. Tao, R. Nair, J. Moreno, N. Muralimanohar, K. Sankaralingam, and C. Estan. Near-Memory Data Services. *IEEE Micro*, 36(1):6–13, Jan 2016.

[35] D. Patterson, T. Anderson, N. Cardwell, R. Fromm, K. Keeton, C. Kozyrakis, R. Thomas, and K. Yelick. Intelligent ram (IRAM): Chips that remember and compute. In *Solid-State Circuits Conference, 1997. Digest of Technical Papers. 43rd ISSCC., 1997 IEEE International*.

[36] M. Oskin, F. T. Chong, and T. Sherwood. Active Pages: A computation model for intelligent memory. In *Computer Architecture, 1998. Proceedings. The 25th Annual International Symposium on*, pages 192–203, Jun 1998.

[37] Erik Riedel, Christos Faloutsos, Garth A Gibson, and David Nagle. Active disks for large-scale data processing. *Computer*, 34(6):68–74, 2001.

[38] Jeff Draper, Jacqueline Chame, Mary Hall, Craig Steele, Tim Barrett, Jeff LaCoss, John Granacki, Jaewook Shin, Chun Chen, Chang Woo Kang, et al. The architecture of the DIVA processing-in-memory chip. In *Proc. ICS*, pages 14–25. ACM, 2002.

[39] R. Nair, S. F. Antao, C. Bertolli, P. Bose, J. R. Brunheroto, T. Chen, C. Y. Cher, C. H. A. Costa, J. Doi, C. Evangelinos, B. M. Fleischer, T. W. Fox, D. S. Gallo, L. Grinberg, J. A. Gunnels, A. C. Jacob, P. Jacob, H. M. Jacobson, T. Karkhanis, C. Kim, J. H. Moreno, J. K. O'Brien, M. Ohmacht, Y. Park, D. A. Prener, B. S. Rosenburg, K. D. Ryu, O. Sallenave, M. J. Serrano, P. D. M. Siegl, K. Sugavanam, and Z. Sura. Active memory cube: A processing-in-memory architecture for exascale systems. *IBM Journal of Research and Development*, 59(2/3):17:1–17:14, March 2015.

[40] D. Kim, J. Kung, S. Chai, S. Yalamanchili, and S. Mukhopadhyay. Neurocube: A Programmable Digital Neuromorphic Architecture with High-Density 3D Memory. In *2016 ACM/IEEE 43rd Annual International Symposium on Computer Architecture (ISCA)*, pages 380–392, June 2016.

[41] A. Farmahini-Farahani, J. H. Ahn, K. Morrow, and N. S. Kim. NDA: Near-DRAM acceleration architecture leveraging commodity DRAM devices and standard memory modules. In *2015 IEEE 21st International Symposium on High Performance Computer Architecture (HPCA)*, pages 283–295, Feb 2015.

[42] S. H. Pugsley, J. Jestes, H. Zhang, R. Balasubramonian, V. Srinivasan, A. Buyuktosunoglu, A. Davis, and F. Li. NDC: Analyzing the impact of 3D-stacked memory+logic devices on MapReduce workloads. In *2014 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, pages 190–200, March 2014.

[43] Dongping Zhang, Nuwan Jayasena, Alexander Lyashevsky, Joseph L. Greathouse, Lifan Xu, and Michael Ignatowski. TOP-PIM: Throughput-oriented Programmable Processing in Memory. In *Proceedings of the 23rd International Symposium on High-performance Parallel and Distributed Computing*, HPDC '14, pages 85–98, New York, NY, USA, 2014. ACM.

[44] J. Ahn, S. Yoo, O. Mutlu, and K. Choi. PIM-enabled instructions: A low-overhead, locality-aware processing-in-memory architecture. In *2015 ACM/IEEE 42nd Annual International Symposium on Computer Architecture (ISCA)*, pages 336–348, June 2015.

[45] Q. Zhu, K. Vaidyanathan, O. Shacham, M. Horowitz, L. Pileggi, and F. Franchetti. Design automation framework for application-specific logic-in-memory blocks. In *2012 IEEE 23rd International Conference on Application-Specific Systems, Architectures and Processors*, pages 125–132, July 2012.

[46] J. T. Pawlowski. Hybrid memory cube (HMC). In *Hot Chips*, volume 23, 2011.

[47] D. Lee et al. 25.2 A 1.2 V 8Gb 8-channel 128GB/s high-bandwidth memory (HBM) stacked DRAM with effective microbump I/O test methods using 29nm process and TSV. In *Proc. ISSCC*, pages 432–433. IEEE, 2014.

[48] T. Hanyu. Challenge of MTJ/MOS-hybrid logic-in-memory architecture for nonvolatile VLSI processor. In *Circuits and Systems (ISCAS), 2013 IEEE International Symposium on*, pages 117–120, May 2013.

[49] M. Natsui, D. Suzuki, N. Sakimura, R. Nebashi, Y. Tsuji, A. Morioka, T. Sugibayashi, S. Miura, H. Honjo, K. Kinoshita, S. Ikeda, T. Endoh, H. Ohno, and T. Hanyu. Nonvolatile Logic-in-Memory LSI using cycle-based power gating and its application to motion-vector prediction. *IEEE Journal of Solid-State Circuits*, 50(2):476–489, Feb 2015.

[50] Shoun Matsunaga, Jun Hayakawa, Shoji Ikeda, Katsuya Miura, Tetsuo Endoh, Hideo Ohno, and Takahiro Hanyu. MTJ-based nonvolatile logic-in-memory circuit, future prospects and issues. In *Proceedings of the Conference on Design, Automation and Test in Europe*, pages 433–435. European Design and Automation Association, 2009.

[51] Meng-Fan Chang, Albert Lee, Chien-Chen Lin, Mon-Shu Ho, Ping-Cheng Chen, Chia-Chen Kuo, Ming-Pin Chen, Pei-Ling Tseng, Tzu-Kun Ku, Chien-Fu Chen, Kai-Shin Li, and Jia-Min Shieh. Read circuits for resistive memory (ReRAM) and memristor-based nonvolatile Logics. In *The 20th Asia and South Pacific Design Automation Conference*, pages 569–574, Jan 2015.

[52] Nios II Processor, Intel Corporation.

[53] D. Lee, X. Fong, and K. Roy. R-MRAM: A ROM-Embedded STT MRAM Cache. *IEEE Electron Device Letters*, 34(10):1256–1258, Oct 2013.

[54] Paul Dlugosch, Dave Brown, Paul Glendenning, Michael Leventhal, and Harold Noyes. An efficient and scalable semiconductor architecture

for parallel automata processing. *IEEE Transactions on Parallel and Distributed Systems*, 25(12):3088–3098, 2014.

[55] Xuanyao Fong, Sri Harsha Choday, Panagopoulos Georgios, Charles Augustine, and Kaushik Roy. Spice models for magnetic tunnel junctions based on monodomain approximation https://nanohub.org/resources/19048, Aug 2016.

[56] S Ikeda, K Miura, H Yamamoto, K Mizunuma, HD Gan, M Endo, Sl Kanai, J Hayakawa, F Matsukura, and H Ohno. A perpendicular-anisotropy CoFeB–MgO magnetic tunnel junction. *Nature materials*, 9(9):721–724, 2010.

[57] Naveen Muralimanohar, Rajeev Balasubramonian, and Norm Jouppi. Optimizing NUCA Organizations and Wiring Alternatives for Large Caches with CACTI 6.0. In *Proceedings of the 40th Annual IEEE/ACM International Symposium on Microarchitecture*.

Shubham Jain is currently a PhD student in the School of Electrical and Computer Engineering, Purdue University. His research interests include exploring circuit and architectural techniques for emerging post-CMOS devices and computing paradigms such as spintronics, approximate computing and neuromorphic computing. He has a B.Tech(Hons.) degree in Electronics and Electrical Communication Engineering from the Indian Institute of Technology, Kharagpur, India, in 2012. After graduation, he worked for two years in Qualcomm, Bangalore, India. He is a recipient of the Andrews Fellowship from Purdue University, in 2014.

Ashish Ranjan received the BTech degree in electronics engineering from the Indian Institute of Technology (BHU), Varanasi, India, in 2009. He is currently working towards a PhD degree in the School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN. His industry experience includes three years as a senior member technical staff in the Design Creation Division, Mentor Graphics Corporation, Noida, India. His primary research interests include circuit-architecture codesign for emerging technologies and approximate computing. He was awarded the University Gold Medal for his academic performance by IIT (BHU), Varanasi in 2009. He also received the Andrews Fellowship from Purdue University in 2012.

Kaushik Roy received the BTech degree in electronics and electrical communications engineering from the Indian Institute of Technology, Kharagpur, India, and the PhD degree from the Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign in 1990. He was with the Semiconductor Process and Design Center of Texas Instruments, Dallas, where he worked on FPGA architecture development and low-power circuit design. He joined the electrical and computer engineering faculty at Purdue University, West Lafayette, IN, in 1993, where he is currently Edward G. Tiedemann Jr. Distinguished Professor. His research interests include spintronics, device-circuit co-design for nano-scale Silicon and non-Silicon technologies, low-power electronics for portable computing and wireless communications, and new computing models enabled by emerging technologies. He has published more than 600 papers in refereed journals and conferences, holds 15 patents, graduated 60 PhD students, and is coauthor of two books on Low Power CMOS VLSI Design (Wiley & McGraw Hill). He received the US National Science Foundation Career Development Award in 1995, IBM faculty partnership award, ATT/Lucent Foundation award, 2005 SRC Technical Excellence Award, SRC Inventors Award, Purdue College of Engineering Research Excellence Award, Humboldt Research Award in 2010, 2010 IEEE Circuits and Systems Society Technical Achievement Award, Distinguished Alumnus Award from Indian Institute of Technology, Kharagpur, Fulbright-Nehru Distinguished Chair, and Best Paper Awards at 1997 International Test Conference, IEEE 2000 International Symposium on Quality of IC Design, 2003 IEEE Latin American Test Workshop, 2003 IEEE Nano, 2004 IEEE International Conference on Computer Design, 2006 IEEE/ACM International Symposium on Low Power Electronics & Design, and 2005 IEEE Circuits and System Society Outstanding Young Author Award (Chris Kim), 2006 IEEE Transactions on VLSI Systems Best Paper Award, 2012 ACM/IEEE International Symposium on Low Power Electronics and Design Best Paper Award, 2013 IEEE Transactions on VLSI Best Paper Award. He was a Purdue University Faculty scholar (1998-2003). He was a Research Visionary board member of Motorola Labs (2002) and held the M.K. Gandhi Distinguished Visiting faculty at Indian Institute of Technology (Bombay). He has been in the editorial board of IEEE Design and Test, IEEE Transactions on Circuits and Systems, IEEE Transactions on VLSI Systems, and IEEE Transactions on Electron Devices. He was the guest editor for Special Issue on Low-Power VLSI in the IEEE Design and Test (1994) and IEEE Transactions on VLSI Systems (June 2000), IEE Proceedings—Computers and Digital Techniques (July 2002), and IEEE Journal on Emerging and Selected Topics in Circuits and Systems (2011). He is a fellow of the IEEE

Anand Ragunathan is a Professor of Electrical and Computer Engineering and Chair of the VLSI area at Purdue University, where he directs research in the Integrated Systems Laboratory. His current areas of research include domain-specific architecture, system-on-chip design, computing with post-CMOS devices, and heterogeneous parallel computing. Previously, he was a Senior Research Staff Member at NEC Laboratories America, where he led projects on system-on-chip architecture and design methodology. He has also held the Gopalakrishnan Visiting Chair in the Department of Computer Science and Engineering at the Indian Institute of Technology, Madras.

Prof. Raghunathan has co-authored a book, eight book chapters, and over 200 refereed journal and conference papers, and holds 21 U.S patents. His publications received eight best paper awards and five best paper nominations. He received a Patent of the Year Award and two Technology Commercialization Awards from NEC, and was chosen among the MIT TR35 (top 35 innovators under 35 years across various disciplines of science and technology) in 2006.

Prof. Raghunathan has been a member of the technical program and organizing committees of several leading conferences and workshops, chaired premier IEEE/ACM conferences (CASES, ISLPED, VTS, and VLSI Design), and served on the editorial boards of various IEEE and ACM journals in his areas of interest. He received the IEEE Meritorious Service Award and Outstanding Service Award. He is a Fellow of the IEEE and Golden Core Member of the IEEE Computer Society. Prof. Raghunathan received the B. Tech. degree from the Indian Institute of Technology, Madras, and the M.A. and Ph.D. degrees from Princeton University.